PassQueen
a.k.a dumpsqueen.com

TMOS Administration

F5 201

Version Demo

Total Demo Questions: 27

Total Premium Questions: 274

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Troubleshoot Basic Virtual Server Connectivity Issues	50
Topic 2, Troubleshoot Basic Hardware Issues	10
Topic 3, Troubleshoot Basic Performance Issues	16
Topic 4, Troubleshoot Basic Device Management Connectivity Issues	10
Topic 5, Troubleshoot Basic Configuration Issues	116
Topic 6, Maintain and Monitor a TMOS Device	72
Total	274

QUESTION NO: 1

Where is connection mirroring configured?

A. It an option within a TCP profile.

B. It is an optional feature of each pool.

C. It is not configured; it is default behavior.

D. It is an optional feature of each virtual server.

ANSWER: D

Explanation:

Connection mirroring is an optional setting configured on an individual virtual server. When enabled, BIG-IP mirrors the connection state for connections handled by that virtual server to its high-availability peer. This state synchronization allows the peer device to continue processing eligible existing connections after a failover, rather than requiring all clients to establish new sessions. The setting is exposed in the virtual server's configuration as the connection-mirroring property, and it can be enabled only where the application's availability requirements justify the associated synchronization overhead. Connection mirroring is commonly used with stateful TCP-based applications when preserving in-progress connections during an active-standby failover is important. It works as part of an overall device-service-clustering configuration, where the devices have the necessary failover and mirroring network connectivity. F5's virtual server API reference identifies connection mirroring as a virtual-server attribute, and F5's connection mirroring documentation describes its role in maintaining connection state across an HA pair. See the F5 BIG-IP virtual server API reference and the F5 CONNECTION::mirror documentation.

QUESTION NO: 2

Assume the bigd daemon fails on the active system. Which three are possible results? (Choose three.)

A. The active system will restart the bigd daemon and continue in active mode.

B. The active system will restart the tmm daemon and continue in active mode.

C. The active system will reboot and the standby system will go into active mode.

D. The active system will failover and the standby system will go into active mode.

E. The active system will continue in active mode but gather member and node state information from the standby system.

ANSWER: A C D

Explanation:

The BIG-IP daemon high-availability framework can take different actions when a monitored daemon such as `bigd` fails, based on the configured `sys daemon-ha` policy and the severity or persistence of the failure. The active system will restart the bigd daemon and continue in active mode is a possible outcome because daemon HA can be configured to restart a failed process locally, restoring health-monitor processing while retaining the current traffic-group role.

The active system will reboot and the standby system will go into active mode is also possible when the configured daemon-HA action is reboot. Rebooting removes the failed device from its active role; in a properly configured redundant pair, the peer can assume the active traffic-group role to maintain service availability. The active system will failover and the standby system will go into active mode is likewise a valid configured response. A failover action transitions the traffic-group role to the peer without requiring the local unit to be rebooted first. These actions allow administrators to select an appropriate recovery response for critical daemon failures. F5 documents the available daemon HA actions, including restart, failover, and reboot, in the tmsh sys daemon-ha reference; health-monitor operation is described in the BIG-IP health monitor documentation.

QUESTION NO: 3

Which two statements describe differences between the active and standby systems? (Choose two.)

A. Monitors are performed only by the active system.

B. Failover triggers only cause changes on the active system.

C. Virtual server addresses are hosted only by the active system.

D. Configuration changes can only be made on the active system.

E. Floating selfIP addresses are hosted only by the active system.

ANSWER: C E

Explanation:

In a high-availability pair configured for active/standby operation, traffic processing is normally centralized on the active device. The active device owns the traffic-group addresses needed to receive client connections and forward traffic, while the standby device is prepared to take ownership if a failover occurs. Therefore, virtual server addresses are hosted only by the active system during normal operation. A virtual server's address is the client-facing destination used for application traffic, so its active traffic-group ownership distinguishes the active unit from its peer.

Floating self IP addresses are likewise traffic-group resources. They move as a group between devices during failover and are hosted by the active system under normal conditions. This floating behavior gives the network a stable address for services such as routing, high availability, and SNAT-related connectivity, without requiring upstream devices to use a different address after the active device changes. When the active device fails or a failover is initiated, the newly active device assumes ownership of these virtual and floating addresses and announces the change to the network. F5 documents traffic groups and their floating-address behavior in its Traffic groups documentation and describes high-availability concepts in the Device Service Clustering overview.

QUESTION NO: 4

Refer to the exhibit.

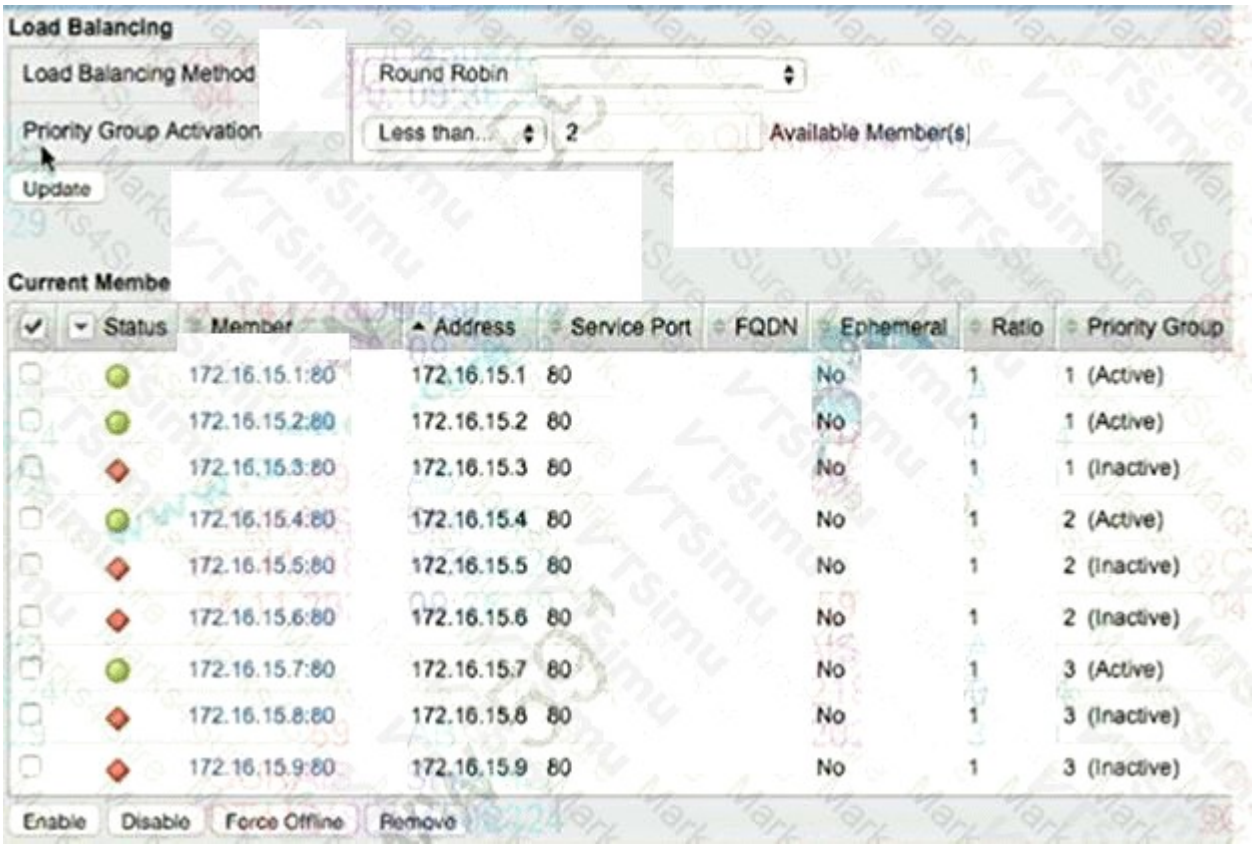

Which two pool members should be chosen for a new connection? (Choose two.)

A. 172.16.15.9.80

B. 172.16.15.4.80

C. 172.10.15.2.80

D. 172.16.15.1.80

E. 172.16.15.7.80

ANSWER: B E

Explanation:

172.16.15.4.80 and **172.16.15.7.80** are the pool members eligible to receive a new connection in the displayed pool status. BIG-IP load balances new flows only to members that are enabled for traffic and considered available by their configured health-monitoring state. The exhibit identifies these two members as active, available destinations, so they participate in the pool's load-balancing decision for a new client-side connection.

Pool-member eligibility is evaluated independently of the pool's overall status. A member must be administratively enabled and have a monitor result that permits traffic before the Local Traffic Manager can select it. When an eligible member is selected, BIG-IP creates the server-side connection to that member's configured address and service port; in this case, the HTTP service on port 80. This behavior prevents traffic from being sent to a member that is not currently usable while allowing healthy enabled members to continue serving requests. F5's documentation describes pool members as the individual traffic destinations and explains that health monitors determine their availability for load balancing. See the F5 Local Traffic Manager pool documentation and the BIG-IP Local Traffic Management Basics guide.

QUESTION NO: 5

A virtual server is listening at 10.10.1.100:80 and has the following iRule associated with it: when HTTP_REQUEST { if {[HTTP::uri] ends_with "txt" } { pool pool1 } elseif {[HTTP::uri] ends_with "php" } { pool pool2 } If a user connects to http://10.10.1.100/foo.html, which pool will receive the request?

A. pool1

B. pool2

C. None. The request will be dropped.

D. Unknown. The pool cannot be determined from the information provided.

ANSWER: D

Explanation:

Unknown. The pool cannot be determined from the information provided.

is correct because the request URI is `/foo.html`, which does not end with either `txt` or `php`. Consequently, neither conditional branch invokes the `pool` iRule command, so this iRule does not explicitly select `pool1` or `pool2` for this request. The `HTTP::uri` command supplies the URI used in the conditional tests, while the `pool` command is the action that directs traffic to a named pool. See the F5 iRules references for HTTP::uri and pool.

Whether the request is ultimately sent to a pool depends on configuration not included in the question, most notably whether the virtual server has a default pool or whether another associated traffic-selection mechanism applies. If a default pool is configured, normal virtual-server processing can use that pool when this iRule does not override the selection. If no applicable pool-selection configuration exists, the connection cannot be load balanced to a pool. Since the virtual server's default-pool and any additional policy or iRule configuration are not stated, no specific receiving pool can be identified.

QUESTION NO: 6

The current status of a given pool member is unknown. Which condition could explain that state?

A. The member has no monitor assigned to it.

B. The member has a monitor assigned to it and the most recent monitor was successful.

C. The member has a monitor assigned to it and the monitor did not succeed during the most recent timeout period.

D. The member's node has a monitor assigned to it and the monitor did not succeed during the most recent timeout period.

ANSWER: A

Explanation:

The member has no monitor assigned to it. is correct because a pool member's availability status is determined by health-monitor results when monitoring is configured. If no health monitor is associated with the pool member—either directly or through the pool—the system has no probe result from which to establish that the member is available or unavailable. It therefore displays the member's monitor/availability state as unknown.

Health monitors actively test a service using a configured method, such as a TCP connection or an HTTP request, at the monitor's specified interval. The BIG-IP system uses the resulting success or failure information to calculate availability. A member can inherit a monitor assigned at the pool level, so an administrator should examine both the member's explicit monitor setting and the effective pool monitor configuration before concluding that no monitor applies. When the effective monitor assignment is truly absent, unknown is an expected and useful indication that no health determination has been made, rather than a confirmation that the application service is healthy. F5's pool and pool-member monitoring documentation describes monitor associations and their role in determining member availability: F5 BIG-IP LTM: Monitoring pools and pool members. See also the tmsh pool reference for monitor configuration details.

QUESTION NO: 7

A site needs to decrypt client HTTPS traffic at the BIG-IP and then establish a new TLS connection to the selected pool member. Which three profile types would normally be associated with the virtual server? (Choose three.)

A. TCP

B. ClientSSL

C. ServerSSL

D. HTTP

E. HTTPS

ANSWER: A B C

Explanation:

TCP, ClientSSL, and ServerSSL are the normal profile types for SSL bridging. The Client SSL profile enables BIG-IP to accept the client-side TLS connection, present the configured certificate, and decrypt the client traffic. The Server SSL profile enables BIG-IP to initiate a separate TLS connection to the selected pool member. A TCP profile provides TCP connection handling for the client-side and server-side flows.

An HTTP profile is not inherently required simply to terminate and re-encrypt TLS traffic. It is required only when BIG-IP must perform HTTP-aware processing, such as HTTP header manipulation, cookie persistence, or HTTP iRules processing. F5 documents client-side and server-side SSL configuration in the BIG-IP SSL profile documentation.

QUESTION NO: 8

Which three virtual server properties participate in matching incoming client traffic to a virtual server? (Choose three.)

A. destination address

B. destination port

C. IP protocol

D. pool member ratio

E. pool loadbalancing mode

F. health monitor receive string

ANSWER: A B C

Explanation:

The destination address, destination port, and IP protocol participate in virtual-server matching. BIG-IP uses these properties to determine which configured virtual server should process an incoming flow. For example, separate virtual servers can use the same destination address with different service ports or IP protocols.

The selected virtual server can then apply its profiles, source address translation settings, iRules, and default pool. The pool members and the pool's load-balancing method are not used to determine which virtual server initially receives the client connection. F5 documents virtual-server destination and protocol properties in the tmsh ltm virtual reference.

QUESTION NO: 9

The BIG-IP Administrator configures an HTTP monitor with a specific receive string. The status is marked

'down'.

Which tool should the administrator use to identify the problem?

A. Ping

B. Health

C. tcpdump

D. ifconfig

ANSWER: C

Explanation:

tcpdump is the appropriate tool because it lets the BIG-IP Administrator capture and inspect the actual HTTP health-monitor exchange between the BIG-IP system and the pool member. An HTTP monitor is considered available only when the target returns a response that satisfies the configured receive string. A packet capture can confirm that the monitor request leaves the applicable BIG-IP self IP, reaches the server, and receives a response. It also exposes the precise response payload, including HTTP headers and body content, so the administrator can determine whether the configured receive string appears exactly as expected. This is particularly useful for identifying practical issues such as an unexpected redirect, a different status response, a virtual-host-dependent response, TLS versus clear-text HTTP traffic, or a response body that does not contain the required text. On BIG-IP, captures are commonly filtered by the pool member address and monitor port; adding verbosity options can show the HTTP content needed to validate the receive string. F5 documents tcpdump as the command-line packet-capture utility for troubleshooting BIG-IP traffic and provides monitor troubleshooting guidance in its knowledge resources. See the F5 tcpdump documentation and the F5 health monitor troubleshooting guidance.

QUESTION NO: 10

Which two can be a part of a pool's definition? (Choose two.)

A. rule(s)

B. profile(s)

C. monitor(s)

D. persistence type

E. loadbalancing mode

ANSWER: C E

Explanation:

A pool definition in BIG-IP LTM includes both health-monitor assignment and a load-balancing method. The **monitor(s)** setting associates one or more health monitors with the pool. BIG-IP uses those monitors to determine the availability of pool members and, depending on the configured monitor rule, whether the pool itself is considered available for traffic. This lets the system avoid selecting members that are not successfully responding to the configured health checks.

The **loadbalancing mode** setting defines the algorithm BIG-IP uses to select an eligible pool member for a new connection. Common choices include Round Robin, Least Connections (member), Ratio (member), and observed- or predictive-method variants. The chosen method is evaluated among pool members that are currently eligible based on availability and other applicable traffic-management settings. These are native pool-level properties in the BIG-IP configuration model and are configured from the pool's Properties and Health Monitors areas or through the corresponding TMOS configuration objects. F5 documents pool health monitors and load-balancing methods as pool configuration settings in its BIG-IP LTM documentation. See F5 BIG-IP LTM pools documentation and F5 pool member availability documentation.

QUESTION NO: 11

A BIG-IP Administrator needs to view the CPU utilization of a particular Virtual Server. Which section of the Configuration Utility should the administrator use for this purpose?

A. Statistics > Module Statistics > Local Traffic > Virtual Addresses

B. Statistics > Module Statistics > Traffic Summary

C. Statistics > Analytics > Process CPU Utilization

D. Statistics > Module Statistics > Local Traffic > Virtual Servers

ANSWER: D

Explanation:

`Statistics > Module Statistics > Local Traffic > Virtual Servers` is the Configuration Utility location used to view statistics for an individual LTM virtual server. This section presents virtual-server-specific operational and traffic statistics, allowing an administrator to select the relevant virtual server and inspect its current and historical performance information in the context of that traffic listener. Because the requirement is scoped to a particular virtual server rather than to the BIG-IP system as a whole or to an individual daemon, the Local Traffic virtual-server statistics view is the appropriate starting point. Virtual servers are the LTM objects that receive client traffic and apply the configured processing and load-balancing behavior, so their statistics page is the relevant per-object monitoring interface. F5 documents the virtual-server statistics command and counters in its tmsh ltm virtual reference. Additional guidance on viewing BIG-IP statistics through the Configuration utility is available in the BIG-IP monitoring and management documentation.

QUESTION NO: 12

Which cookie persistence method requires the fewest configuration changes on the web servers to be implemented correctly?

A. insert

B. rewrite

C. passive

D. session

ANSWER: A

Explanation:

The **insert** cookie persistence method requires the fewest web-server configuration changes because BIG-IP creates and inserts its own persistence cookie into the HTTP response sent to the client. On later requests, BIG-IP reads that cookie and uses its encoded information to direct the client back to the previously selected pool member. The application and web servers do not need to generate, recognize, process, or modify the persistence cookie; persistence is handled entirely by the BIG-IP system and its cookie persistence profile.

This makes insert persistence especially useful when an application does not already issue a suitable session cookie, when modifying application code is impractical, or when administrators need a straightforward persistence implementation without coordinating changes across multiple web servers. Configuration is primarily limited to creating or selecting a cookie persistence profile with the insert method and associating it with the relevant virtual server. BIG-IP can use its default cookie name or an administrator-defined name, while maintaining the persistence mapping in its own session table. F5 describes cookie persistence configuration and the insert method in its Local Traffic Manager implementation documentation. Additional persistence-profile concepts are available in the F5 HTTP cookie reference.

QUESTION NO: 13

Which statement is true concerning the default communication between a redundant pair of BIG-IP devices?

A. Communication between the systems cannot be effected by port lockdown settings.

B. Data for both connection and persistence mirroring are shared through the same TCP connection.

C. Regardless of the configuration, some data is communicated between the systems at regular intervals.

D. Connection mirroring data is shared through the serial fail over cable unless network failover is enabled.

ANSWER: B

Explanation:

Data for both connection and persistence mirroring are shared through the same TCP connection. In a redundant BIG-IP device pair, mirroring replicates runtime state from the active device to its peer so that, if a failover occurs, the newly active device has the state needed to continue handling applicable traffic with minimal disruption. This mirrored state can include connection information and persistence-session information, depending on the virtual server and profile configuration. BIG-IP uses its established mirroring communication channel between the peers to convey these state updates; the communication is TCP-based to provide reliable, ordered delivery of the mirrored records. This is distinct from configuration synchronization, which distributes saved configuration objects, and from failover communication, which determines device availability and active/standby state. The shared mirroring channel is therefore a key part of maintaining traffic-processing continuity during a failover, particularly for applications that rely on persistence or long-lived connections. F5 describes connection mirroring and persistence mirroring as high-availability features in its Device Service Clustering documentation. See also F5's network ports used by BIG-IP reference for inter-device communication details.

QUESTION NO: 14

A BIG-IP Administrator needs to collect HTTP status code and HTTP method for traffic flowing through a

virtual server.

Which default profile provides this information?

A. HTTP

B. Analytics

C. Request Adapt

D. Statistics

ANSWER: D

Explanation:

The **Statistics** profile is the appropriate default profile because it enables BIG-IP to collect HTTP traffic statistics associated with a virtual server. Its HTTP statistics include request-method information and response status-code information, allowing an administrator to see the methods clients use, such as GET or POST, and the status codes returned by the application, such as 200, 404, or 500. This is useful for operational monitoring, troubleshooting application behavior, and identifying traffic patterns without modifying the application itself. The profile is attached to the virtual server so that BIG-IP can observe and record these details as it processes matching HTTP flows. The Statistics profile is specifically intended for gathering this type of protocol-level traffic data and can be used in conjunction with BIG-IP reporting and monitoring capabilities. F5 documents the available configuration and collected-statistics behavior for this profile in the tmsh statistics-profile reference. For additional context on applying profiles to virtual servers, see the BIG-IP Local Traffic Manager virtual server documentation.

QUESTION NO: 15

Which two statements describe characteristics of a OneConnect profile? (Choose two.)

A. It can allow multiple eligible client requests to reuse a server-side connection.

B. Its source mask controls client-address matching for connection reuse.

C. It encrypts traffic between BIG-IP and the pool members.

D. It performs source address translation for client traffic.

E. It replaces the pool loadbalancing method.

ANSWER: A B

Explanation:

A OneConnect profile enables BIG-IP to reuse an established server-side connection for subsequent eligible client-side requests. This can reduce the number of server-side TCP connections required when multiple client requests can be multiplexed onto reusable connections.

The source mask setting controls which client source addresses are eligible to share a server-side connection. A mask of `0.0.0.0`, for example, permits reuse without requiring matching client source addresses. OneConnect is commonly used with HTTP traffic so that BIG-IP can identify individual requests. F5 documents OneConnect behavior and source-mask configuration in the F5 tmsh OneConnect profile reference.

QUESTION NO: 16

Which three methods can be used for initial access to a BIG-IP system? (Choose three.)

A. CLI access to the serial console port

B. SSH access to the management port

C. SSH access to any of the switch ports

D. HTTP access to the management port

E. HTTP access to any of the switch ports

F. HTTPS access to the management port

G. HTTPS access to any of the switch ports

ANSWER: A B F

Explanation:

Initial access to a BIG-IP system is available through its local console connection or through the dedicated management interface. “CLI access to the serial console port” provides direct, out-of-band terminal access to the appliance or virtual-edition console. This method is particularly important for first-time setup, because it does not depend on a previously reachable network management address and enables an administrator to configure the initial management settings.

After the management interface has an IP address and is reachable, “SSH access to the management port” provides secure command-line administration. Administrators can authenticate and use the BIG-IP command-line environment to inspect or configure the system. “HTTPS access to the management port” provides secure browser-based access to the Configuration utility, including the setup utility used to complete initial configuration tasks. F5 documents use of the console for initial platform configuration and identify the management interface as the administrative access path for secure SSH and HTTPS connections. See the F5 BIG-IP System: Getting Started documentation and F5 guidance on the default management IP address.

QUESTION NO: 17

A BIG-IP Administrator plans to upgrade a BIG-IP device to the latest TMOS version.

Which two tools could the administrator leverage to verify known issues for the target versions?

(Choose two.)

A. F5 University

B. F5 Downloads

C. F5 End User Diagnostics (EUD)

D. F5 iHealth

E. F5 Bug Tracker

ANSWER: B E

Explanation:

F5 Downloads and F5 Bug Tracker are the appropriate resources for assessing known issues before selecting and installing a target TMOS release. F5 Downloads provides access to the target BIG-IP software image and, importantly, the associated release documentation. BIG-IP release notes contain a documented known-issues section for that release, along with upgrade considerations, behavioral changes, and applicable conditions or workarounds. Reviewing these notes lets an administrator determine whether a reported condition applies to the modules and configuration in the planned deployment.

F5 Bug Tracker provides the more granular, continuously maintained view of individual product defects. An administrator can search or filter bugs by BIG-IP version, product area, and status to identify defects affecting the intended target version and to review symptoms, affected configurations, and available workaround information. Used together, the release notes from F5 Downloads provide the release-level upgrade review, while F5 Bug Tracker supports detailed validation of specific known defects before the maintenance window. See the F5 Downloads portal and the F5 Bug Tracker.

QUESTION NO: 18

A virtual server is configured with VLAN Traffic set to Enabled on VLANs and includes only vlan_clients. Which two statements are true? (Choose two.)

A. Traffic arriving on vlan_clients can match the virtual server.

B. Traffic arriving on a VLAN not in the enabled list cannot match the virtual server.

C. The virtual server accepts traffic on every VLAN by default.

D. The VLAN restriction applies only to server-side traffic.

E. The virtual server can no longer use a pool.

ANSWER: A B

Explanation:

With VLAN Traffic set to Enabled on, the virtual server accepts traffic only when it arrives on a VLAN included in its enabled VLAN list. Therefore, a client connection arriving on vlan_clients is eligible to match the virtual server if the other destination and protocol criteria also match.

A connection arriving on a VLAN that is not included in the enabled list does not match the virtual server because of that virtual server's VLAN restriction. VLAN restrictions are useful when the same virtual address must be exposed only to selected network segments. F5 documents virtual-server VLAN controls in the F5 tmsh ltm virtual reference.

QUESTION NO: 19

A BIG-IP Administrator creates an HTTP Virtual Server using an iApp template. After the Virtual Server is

created, the user requests to change the destination IP addresses. The BIG-IP Administrator tries to

change the destination IP address from 10.1.1.1 to 10.2.1.1 in Virtual Server settings, but receives the

following error:

The application service must be updated using an application management interface

What is causing this error?

A. The Application Service was NOT deleted before making the IP address change.

B. The IP addresses are already in use.

C. The Application Services have Strict Updates enabled.

D. The IP addresses used are NOT from the same subnet as the Self IP.

ANSWER: C

Explanation:

The Application Services have Strict Updates enabled. An iApp template creates and manages the virtual server and its related BIG-IP objects as members of an application service. When Strict Updates is enabled for that application service, TMOS prevents administrators from directly modifying managed objects through their individual configuration screens or through equivalent direct object-level changes. This preserves the intended template-driven configuration and prevents the application service from drifting from the settings defined through the iApp management workflow.

Because the virtual server destination address is a property managed by the deployed iApp application service, the administrator must update the value through the application service's management interface, such as the iApps/Application Services interface, rather than editing the virtual server directly. The displayed error is the expected protection mechanism for a strictly managed application service. F5 documents that strict updates restrict direct modification of application-service objects and require updates through the application service itself. See the F5 Application Service documentation and the F5 iApps documentation.

QUESTION NO: 20

A BIG-IP Administrator needs to modify a virtual server that web offload web traffic compression tasks from the target server.

Which two profiles must the BIG-IP Administrator apply to a virtual server to enable compression? (Choose two)

A. Server SSL profile

B. Stream profile

C. Persistence profile

D. HTTP profile

E. Compression profile

ANSWER: D E

Explanation:

An HTTP profile and a Compression profile must be associated with the virtual server to perform HTTP response compression on BIG-IP. The HTTP profile enables BIG-IP to interpret and process HTTP traffic, including HTTP headers that govern compression negotiation. In particular, BIG-IP uses the client's `Accept-Encoding` header to determine whether the client supports a compression method and can receive a compressed response. The Compression profile supplies the compression behavior and policy, such as the compression algorithms, content types, URI exclusions, and minimum response size to which compression applies. Together, these profiles allow BIG-IP to compress eligible server responses before delivering them to clients, reducing response payload sizes and shifting that processing work away from the target web servers. For HTTPS traffic, SSL decryption must also be configured as appropriate so BIG-IP can inspect the HTTP content, but an SSL profile is not intrinsically one of the two profiles that enables HTTP compression. F5 documents HTTP profile settings through its HTTP profile API reference and compression-profile configuration through its Compression profile API reference.

QUESTION NO: 21

Which statement is true about the synchronization process, as performed by the Configuration Utility or by typing b config sync all?

A. The process should always be run from the standby system.

B. The process should always be run from the system with the latest configuration.

C. The two /config/bigip.conf configuration files are synchronized (made identical) each time the process is run.

D. Multiple files, including /config/bigip.conf and /config/bigip_base.conf, are synchronized (made identical) each time the process is run.

ANSWER: C

Explanation:

The two /config/bigip.conf configuration files are synchronized (made identical) each time the process is run. In this legacy BIG-IP ConfigSync workflow, `b config sync all` synchronizes the shared BIG-IP traffic-management configuration represented by `/config/bigip.conf` between the paired systems. The synchronization operation ensures that the peer receives the same object definitions and settings contained in that configuration file, such as virtual servers, pools, pool members, profiles, monitors, and related Local Traffic Manager objects. Consequently, after a successful synchronization, the participating devices have matching `bigip.conf` contents for the configuration data that ConfigSync manages.

This distinction matters in a high-availability pair because the configuration must be consistent on both units for failover to preserve the intended application-delivery behavior. F5 documents that the BIG-IP configuration is stored in configuration files under `/config`, with `bigip.conf` holding the primary traffic-management configuration, and describes ConfigSync as the mechanism used to synchronize configuration data between devices. See the F5 ConfigSync documentation and the F5 configuration-files overview.

QUESTION NO: 22

A BIG-IP system must reach an NTP server through its dedicated management interface. Which type of route should be configured for this purpose?

A. A management route

B. A default pool route

C. A virtual address route

D. A route domain route

ANSWER: A

Explanation:

A management route is correct because it directs management-plane traffic through the dedicated management interface. Management services such as NTP, DNS, SSH administration, and system updates can use management routes without relying on the TMM data-plane routing table.

Standard routes are used for traffic processed through TMM interfaces, VLANs, and self IP addresses. A management route is maintained separately and is appropriate when the destination must be reached through the management network. F5 documents management routes in the tmsh sys management-route reference.

QUESTION NO: 23

A BIGJP Administrator needs to load a UCS file but must exclude the license file. How should the administrator perform this task?

A. From the CLI with command U tmsh load /$ys ucs no-license

B. From the GUI, select the UCS file, unchcck the license box, and click restore

C. From the CLI with command tmsh load sys ucs no-license

D. From the GUI, select the UCS file and click restore

ANSWER: C

Explanation:

The appropriate method is to use the TMOS command `tmsh load sys ucs <ucs filename> no-license`. A UCS archive can contain system configuration, certificates and keys, platform-related files, and the BIG-IP license. The `no-license` parameter instructs TMOS to restore the UCS contents while deliberately omitting the license file. This is especially important when moving a configuration to a different BIG-IP device or appliance, because each system should retain the license that is already installed and activated for that specific unit.

The command is run from the TMOS shell and uses the `sys ucs` command hierarchy. Replace `<ucs filename>` with the UCS archive name or applicable path available on the BIG-IP system. Before loading a UCS, administrators should ensure the archive is compatible with the target BIG-IP version and platform, and should maintain a current backup of the target configuration. F5 documents the UCS load operation and its `no-license` option in the TMOS command reference: F5 TMOS sys ucs command reference. Additional UCS backup and restore guidance is available in the F5 UCS archive documentation.

QUESTION NO: 24

A BIG-IP Administrator is receiving intermittent reports from users that SSL connections to the BIG-IP device are failing. Upon checking the log files, the BIG-IP Administrator notices the following error message:

ere tmm < instance > [< pid >]: 01260008:3: SSL transaction (TPS) rate limit reached

After reviewing statistics, the BIG-IP Administrator notices there are a maximum of 1200 client-side SSL

TPS and a maximum of 800 server-side SSL TPS.

What is the minimum SSL license limit capacity the BIG-IP Administrator should upgrade to handle this

peak?

A. 2000

B. 400

C. 800

D. 1200

ANSWER: A

Explanation:

2000 is the minimum required SSL TPS license capacity. BIG-IP SSL TPS licensing accounts for SSL handshakes processed by the Traffic Management Microkernel (TMM). Client-side and server-side SSL handshakes both consume capacity from the SSL transaction rate limit. At the observed peak, BIG-IP processes up to 1200 client-side SSL transactions per second and 800 server-side SSL transactions per second. The required licensed capacity is therefore the aggregate peak rate: 1200 + 800 = 2000 SSL TPS.

This calculation is especially important for virtual servers that terminate client SSL and subsequently establish a separate SSL connection to pool members using a server SSL profile. In that full-proxy SSL deployment, one end-user connection can require distinct client-side and server-side SSL processing, and each handshake contributes to the licensed transaction rate. Selecting a 2000 TPS license accommodates the reported peak and prevents the TMM SSL transaction rate-limit message from occurring under the same workload. Administrators should also allow reasonable headroom above observed peak traffic when selecting a production license tier, because short traffic bursts, certificate renegotiation behavior, and traffic

growth can increase the actual SSL handshake rate. See the F5 SSL traffic-management documentation and the F5 BIG-IP licensing documentation.

QUESTION NO: 25

Which two settings can be configured on a virtual address? (Choose two.)

A. ARP

B. ICMP echo

C. loadbalancing mode

D. health monitor

E. persistence type

ANSWER: A B

Explanation:

A virtual address has ARP and ICMP echo settings. The ARP setting controls whether BIG-IP responds to ARP requests for the virtual address on the relevant VLAN. The ICMP echo setting controls whether BIG-IP responds to ICMP echo requests sent to that virtual address.

These settings apply to the virtual address object, which can be shared by one or more virtual servers. Pool health monitors and load-balancing methods are configured on pools rather than on virtual addresses. F5 documents virtual-address properties in the tmsh ltm virtual-address reference.

QUESTION NO: 26

Refer to the exhibit.

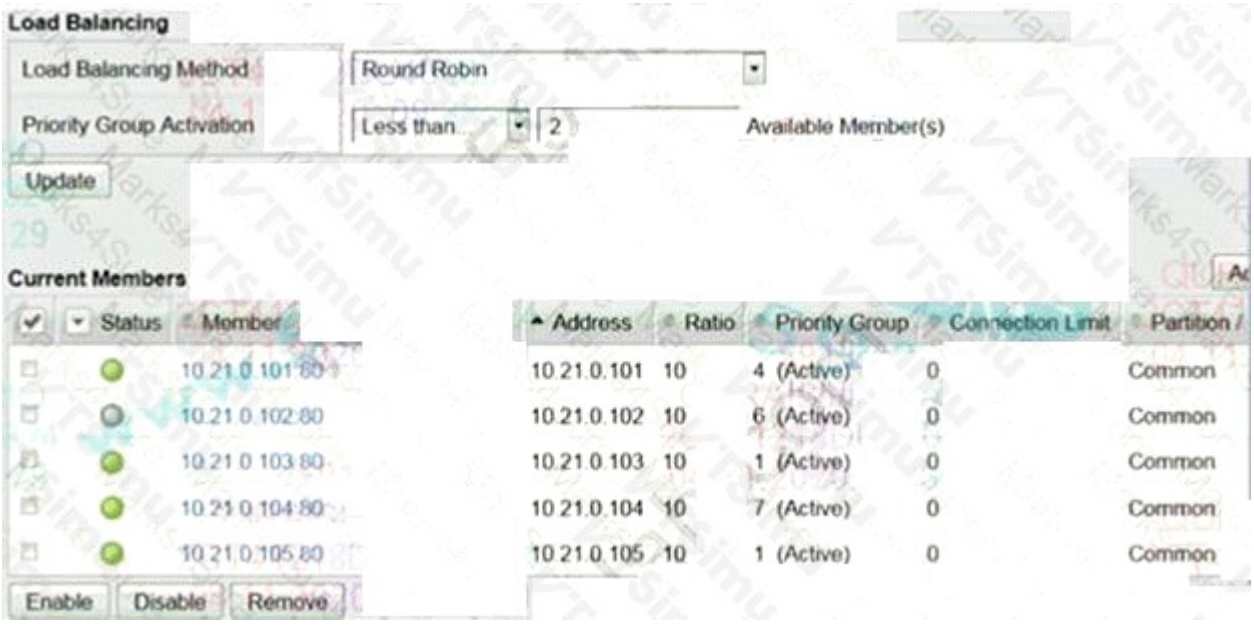

Which two pool members are eligible to receive new connections? (Choose two)

A. 10.21.0.102.80

B. 10.21.0.104.80

C. 10.21.0.105.80

D. 10.21.0.101.80

E. 10.21.0.103.80

ANSWER: B D

Explanation:

10.21.0.104.80 and **10.21.0.101.80** are eligible because the exhibit identifies these pool members as available for traffic. In BIG-IP Local Traffic Manager, a member must be enabled administratively and have an acceptable availability status, normally established by its configured health monitor, before LTM selects it for a new connection. An available member is considered usable by the pool's load-balancing algorithm, subject to any applicable connection limits, priority-group activation, and persistence rules.

The member availability display is therefore the key information for this question: both listed members meet the operational conditions required for normal new-connection selection. BIG-IP continuously evaluates monitored pool members and updates their status so that load balancing is directed only to members that are presently eligible. This behavior provides the core high-availability function of an LTM pool: healthy, enabled members can accept new client flows while the pool maintains service delivery according to its configured traffic-management policy.

See F5's Monitoring pool member status documentation and the About pool members documentation.

QUESTION NO: 27

What is required for a virtual server to support clients whose traffic arrives on the internal VLAN and pool members whose traffic arrives on the external VLAN?

A. That support is never available.

B. The virtual server must be enabled for both VLANs.

C. The virtual server must be enabled on the internal VLAN.

D. The virtual server must be enabled on the external VLAN.

ANSWER: C

Explanation:

The virtual server must be enabled on the internal VLAN. A virtual server's VLAN and Tunnel Traffic settings control which VLANs are permitted to deliver client-side traffic to that virtual server. Since the clients initiate connections from the internal VLAN, that VLAN must be included in the virtual server's enabled VLAN list (or VLAN restrictions must otherwise permit it). This association applies to the packet arriving at the virtual server's destination address and port—the client-facing, ingress side of the connection.

After BIG-IP accepts the client connection and selects a pool member, it routes the server-side flow according to its routing and VLAN configuration. The VLAN used to reach the selected pool member is therefore determined by the applicable route, traffic group/addressing, and egress network connectivity; it does not need to be enabled as an ingress VLAN on the virtual server merely because pool-member traffic uses it. This configuration supports an internal client-side VLAN and an external server-side VLAN while maintaining the intended traffic direction. F5 documents virtual-server VLAN restrictions and their effect on traffic acceptance in the Virtual Servers documentation and describes forwarding and routing behavior in the TMOS Routing Administration guide.